ACTIAN™
a division of HCLSoftware

WHITE PAPER

AI Isn't the Problem. Your Data Model Is.

Why the next analytics revolution will be about meaning, not models

By Emma McGrattan
CTO, Actian

Table of Contents

Executive Summary

Artificial intelligence has never been more capable. Large language models can summarize research papers, write software, and pass professional exams. Yet inside many enterprises, AI systems still struggle to answer basic business questions.

Ask a seemingly simple question such as "How and why did the purchase price fluctuate last quarter?" and the system often produces inconsistent or incomplete answers. The immediate reaction is usually to blame the model. Perhaps the organization needs a more advanced AI system. Perhaps prompts need refinement. Perhaps the training data is insufficient.

In most cases, none of those explanations is correct. The real problem is the data itself.

Enterprise data environments were never designed for AI systems that interpret business questions directly. They were designed for dashboards and reporting workflows where human analysts translate questions into queries. For years, that translation layer hid many structural problems in the data itself. AI removes that layer.

When a system attempts to answer business questions directly, it encounters a landscape filled with inconsistent definitions, fragmented data models, and missing context. The result is not intelligent automation. It is automated confusion.

This is why many organizations are discovering a gap between the promise of AI analytics and the reality of deploying it against enterprise data. The models are powerful. The data environment is not prepared to support them.

The next generation of analytics will not be defined by better models, but by better meaning and the ability to make that meaning accessible to both humans and machines.

About The Author

Emma McGrattan
CTO, Actian

Emma McGrattan is the Chief Technology Officer at Actian, where she leads the strategy behind some of the industry's most innovative data solutions. With over two decades of experience delivering mission-critical data technologies, she has helped organizations across industries design and modernize their data architectures, embed governance by design, and become AI-ready.

ACTIAN™
a division of HCLSoftware

"At Actian, we see a consistent pattern across industries and organizations adopting AI. The ones making the most meaningful progress are not those chasing the latest model releases. They are the ones doing the harder, less glamorous work of making their data environments semantically coherent. They are defining what their metrics mean, establishing clear data ownership, and building the infrastructure for machines to reason reliably rather than plausibly. That investment is what separates organizations that will thrive in the next era of analytics from those that will remain stuck explaining why their AI "almost" works."

– Emma McGrattan, CTO, Actian

The Analytics Capacity Problem

Most organizations are experiencing an explosion in demand for data-driven insight. Executives want faster answers. Product teams want deeper analysis. Operational teams want predictive guidance rather than retrospective reporting. At the same time, the supply of analytical capacity has not kept pace.

For decades, the operating model for analytics has been built around human intermediaries. Business teams ask questions. Analysts interpret those questions, locate the relevant data, reconcile definitions, and construct the necessary queries. The output typically appears in the form of a dashboard, report, or visualization.

This model worked when analytical demand was limited, and the number of data sources was manageable. It begins to break down when every decision is expected to be data-driven, and every team expects immediate answers.

Hiring more analysts does not fundamentally solve the problem. The bottleneck is not simply the number of people performing analysis. The bottleneck is the complexity involved in interpreting enterprise data.

Hiring more analysts does not fundamentally solve the problem. The bottleneck is not simply the number of people performing analysis. The bottleneck is the complexity involved in interpreting enterprise data.

Every analysis requires analysts to answer questions such as:

- **What does this metric actually mean?**
- **Which dataset represents the authoritative version of this information?**
- **How should this metric be calculated across different regions or product lines?**
- **Which transformations have already been applied to the data?**

These decisions are often resolved through experience, institutional knowledge, and informal documentation rather than through explicit structure in the data environment itself.

Human analysts quietly perform this interpretation step every day. Most organizations underestimate how much intellectual work is happening behind the scenes before a chart ever appears on a dashboard. AI systems do not possess that context. When AI attempts to answer questions directly, it encounters ambiguity that humans have been compensating for manually.

Dashboards Were Never Designed to Answer Questions

Dashboards have been the dominant interface for analytics for more than two decades. They are useful tools for monitoring the health of a business. They make it easy to observe trends and identify changes in key performance indicators, but dashboards were never designed to answer open-ended questions.

A dashboard might tell you that revenue declined in a particular region. It will not explain why. It might reveal that churn increased among a certain group of customers. It will not determine which underlying factors contributed to that change. That gap shifts the burden of interpretation back to the user, who must recognize that the dashboard cannot answer their question, formulate a more specific request, and then route it to an analyst; reintroducing the same human intermediary the dashboard was supposed to reduce reliance on.

ACTIAN™ a division of HCLSoftware

When new questions arise, analysts often build new dashboards. Over time, this leads to environments where hundreds or even thousands of dashboards coexist. Each one represents a frozen interpretation of a particular question at a particular moment in time.

This approach introduces several limitations:

- **Dashboards embed business logic inside visual artifacts:** The meaning of a metric becomes tied to a particular chart or report rather than being defined consistently across the organization.
- **Dashboards multiply rapidly:** Each new analytical question produces another dashboard, which introduces additional definitions and interpretations.
- **Dashboards separate insights from context:** Users see numbers and trends without necessarily understanding how those numbers were derived.

AI changes expectations. Instead of navigating dozens of dashboards, business users increasingly expect to ask a question directly and receive a meaningful answer. That shift requires data environments that understand business concepts in a structured and consistent way.

Most current architectures were never designed for that purpose.

AI Will Not Replace BI.
It Will Expose Its Limits

There is a growing debate about whether AI will replace business intelligence. In practice, the relationship is more nuanced.

Dashboards will not disappear. They remain effective tools for monitoring known metrics and tracking performance over time. They provide stability and shared visibility across the organization. However, they are not designed to handle exploratory, open-ended questions.

AI introduces a new interaction model. Instead of navigating predefined views, users can ask questions dynamically and expect contextual answers. The result is not replacement, but convergence.

Organizations will rely on both:

- **BI systems** for monitoring and consistency
- **AI-driven systems** for exploration and reasoning

AI does not eliminate BI. It exposes its limitations and extends what analytics can do. The organizations that understand this distinction will avoid a costly mistake: treating AI as a replacement for BI infrastructure, only to discover the two serve fundamentally different purposes. BI answers known questions reliably. AI explores unknown ones, but only when the data environment gives it the meaning it needs to do so.

The Hidden Problem: Data Without Meaning

Most enterprise data models describe structure rather than meaning. This matters because when meaning is absent from the data layer, interpretation must happen at query time, and that is precisely where AI systems are most vulnerable. A model asked to answer a business question without encoded meaning has no choice but to guess definitions, infer relationships, and construct context on the fly. The result is answers that may be fluent and confident while being factually wrong.

Tables define columns and data types. Pipelines move data from one system to another. Schemas describe relationships between datasets. These structures are necessary for storing and processing data efficiently, but they rarely encode the business meaning of the data.

Consider a simple metric such as customer churn. At first glance, the definition may appear straightforward. In practice, it can vary widely depending on the context. Some organizations measure churn based on account cancellations. Others measure it based on revenue loss or inactivity over a defined period. Different teams often use slightly different definitions depending on their goals.

These differences may seem small, but they accumulate across hundreds of metrics and thousands of datasets. Over time, the result is a data environment where definitions are inconsistent, and meaning must be interpreted rather than retrieved.

ACTIAN™
a division of HCLSoftware

Human analysts can navigate this complexity because they understand the organization and its business context. Machines cannot rely on intuition. AI systems require data environments where meaning is explicit rather than implied. Without that structure, the system must guess which definitions apply and how datasets relate to one another. The answers it produces may appear plausible while still being incorrect.

This is one reason many AI analytics demonstrations appear impressive in controlled environments but struggle when deployed in production. The model is capable of reasoning. The data lacks shared meaning.

The Cost of Ambiguity

There is also a challenge that is not technical, but economic. The cost of AI is often framed in terms of infrastructure: compute, storage, and model training. But one of the highest and least visible costs is ambiguity in the data itself.

Every inconsistent definition, undocumented transformation, or unclear relationship introduces friction:

- AI systems generate plausible but incorrect answers
- Analysts must validate and rework results
- Business users lose trust in automated insights

In many organizations, the cost of validating AI-generated answers now rivals, or even exceeds, the cost of generating them. These inefficiencies compound over time. Poorly structured data environments do not just limit AI performance. They make AI expensive to operate at scale.

For finance leaders evaluating AI investments, this is a critical consideration. The ROI case for AI analytics erodes quickly when validation overhead consumes analyst time that was supposed to be freed up. Semantic data quality is not a technical nice-to-have — it is a direct determinant of whether AI delivers measurable financial return or simply shifts costs from one part of the organization to another.

Why AI Exposes Weaknesses in the Data Stack

AI does not create new data problems. It shines a spotlight on the ones that have been present all along.

Traditional analytics workflows contained a buffer between business users and raw data. Analysts interpreted requests, reconciled inconsistencies, and applied the appropriate transformations before presenting results. That process masked many structural issues in enterprise data environments.

When AI attempts to perform analysis directly, those issues become visible. The system encounters multiple definitions of the same metric. It discovers datasets with unclear ownership. It identifies pipelines where transformations are poorly documented or inconsistently applied. Each of these issues increases the likelihood that the system will misinterpret the data.

This is why many organizations report a gap between successful AI prototypes and reliable production deployments. Early demonstrations often rely on carefully curated datasets where definitions are clear, and relationships are well understood. Enterprise environments are rarely that clean.

Deploying AI at scale requires addressing the underlying data architecture rather than focusing solely on model performance.

The Next Phase of Analytics

The next phase of analytics will focus less on pipelines and visualizations and more on meaning.

Data environments must evolve so that business concepts are defined explicitly and consistently. Instead of relying on dashboards to encode definitions, organizations will need structures that describe how metrics are calculated, how datasets relate to one another, and how context should be interpreted.

Metadata will play a more active role in this environment. Information about data ownership, lineage, quality, and usage will become essential for both governance and machine interpretation.

At the same time, organizations are beginning to treat datasets as governed data products rather than raw assets. These products include clear definitions, quality expectations, and usage constraints so that downstream systems can rely on them with confidence.

These changes are not simply technical improvements. They represent a shift in how organizations think about data. **Data is no longer just a collection of tables waiting to be queried. It becomes a structured representation of business knowledge.**

From Queries to Questions

The analytics systems of the past were designed around queries. The analytics systems of the future will be designed around questions.

When a business leader asks why churn increased in a specific region, the system must understand:

- What churn means
- Which datasets define it
- How it should be calculated
- Which contextual relationships influence the result.

Answering that question requires more than retrieving numbers from a database. It requires an understanding of how the business operates. Organizations that continue to treat data primarily as structured storage will struggle to deliver this capability.

Organizations that ignore this will find themselves on a different trajectory: more models, more pipelines, more dashboards, and still no reliable answers. They will scale the appearance of AI adoption without scaling its benefits, accumulating technical debt in the form of unresolved semantic ambiguity. Those who invest in capturing meaning within their data environments will unlock a different analytical experience. Business users will be able to explore questions directly, while AI systems will interpret data with the same contextual awareness that analysts apply today.

The Emergence of AI Analysts

This shift is also changing the role of analytics systems themselves. Instead of serving as passive reporting tools, they are becoming active participants in decision-making and capable of interpreting questions, retrieving context, and generating insights dynamically.

This is the foundation of a new class of systems often described as AI analysts: systems that augment or automate parts of the analytical process.

However, these systems can only operate effectively if the underlying data environment provides consistent meaning and context. Without that foundation, AI does not scale analysis. It scales ambiguity.

Actian's AI Analyst is built on exactly this premise. Rather than layering a conversational interface on top of raw data, it operates against a semantically governed data environment where business definitions are explicit, metrics are consistently calculated, and context is encoded rather than assumed. The result is an AI analyst that can be trusted to answer business questions with the same rigor a skilled human analyst would apply at a fraction of the time and cost.

Takeaway

Artificial intelligence will continue to improve. Models will become faster, more accurate, and more capable.

But the organizations that benefit most from those advances will not necessarily be the ones with the most advanced models. They will be the ones whose data environments allow machines to understand what the data actually means.

The next generation of analytics will not be defined by better models, but by better meaning. Organizations that continue to optimize dashboards and pipelines without addressing semantics will struggle to scale AI.

Those who invest in making their data understandable to both humans and machines will unlock a fundamentally different analytical experience.

About Actian

Actian empowers enterprises to confidently manage and govern data at scale. Organizations trust Actian data management and data intelligence solutions to streamline complex data environments and accelerate the delivery of AI-ready data. Designed to be flexible, Actian solutions integrate seamlessly and perform reliably across on-premises, cloud, and hybrid environments. Learn more about Actian, the data and AI division of HCLSoftware, at actian.com.

Tel +1 512 231 6000 Fax +1.512.231.6010 710 Hesters Crossing Road, Suite 250, Round Rock, TX 78681 **Actian.com**

ACTIAN™ a division of **HCLSoftware**